arXiv:1604.02264v1 [cs.LG] 8 Apr 2016

Probabilistic classifiers with low rank indefinite kernels

Frank-Michael Schleif[1]*, Andrej Gisbrecht[2], Peter Tino[1]

[1] *School of Computer Science, University of Birmingham, Birmingham B15 2TT, UK*

[2] *Helsinki Institute for Information Technology, Department of Computer Science, Aalto University, Finland*

Abstract

Indefinite similarity measures can be frequently found in bio-informatics by means of alignment scores, but are also common in other fields like shape measures in image retrieval. Lacking an underlying vector space, the data are given as pairwise similarities only. The few algorithms available for such data do not scale to larger datasets. Focusing on probabilistic batch classifiers, the Indefinite Kernel Fisher Discriminant (iKFD) and the Probabilistic Classification Vector Machine (PCVM) are both effective algorithms for this type of data but, with cubic complexity. Here we propose an extension of iKFD and PCVM such that linear runtime and memory complexity is achieved for low rank indefinite kernels. Employing the Nyström approximation for indefinite kernels, we also propose a new almost parameter free approach to identify the landmarks, restricted to a supervised learning problem. Evaluations at several larger similarity data from various domains show that the proposed methods provides similar generalization capabilities while being easier to parametrize and substantially faster for large scale data.

Keywords: indefinite kernel, kernel fisher discriminant, minimum enclosing ball, Nyström approximation, low rank approximation, classification

1. Introduction

Domain specific proximity measures, like alignment scores in bioinformatics Smith et al. (1981), the modified Hausdorff-distance for structural pattern recog-

*Corresponding author

Email addresses: schleify@cs.bham.ac.uk (Frank-Michael Schleif[1]), andrej.gisbrecht@aalto.fi (Andrej Gisbrecht[2]), pxt@cs.bham.ac.uk (Peter Tino[1])

nition Dubuisson and Jain (1994), shape retrieval measures like the inner distance Ling and Jacobs (2007) and many other ones generate non-metric or indefinite similarities or dissimilarities. Classical learning algorithms like kernel machines assume Euclidean metric properties in the underlying data space and may not be applicable for this type of data.

Only few machine learning methods have been proposed for non-metric proximity data, like the indefinite kernel Fisher discriminant (iKFD) Haasdonk and Pekalska (2008); Pekalska and Haasdonk (2009), the probabilistic classification vector machine (PCVM) Chen et al. (2009a) or the indefinite Support Vector Machine (iSVM) in different formulations Haasdonk (2005); Alabdulmohsin et al. (2014); Loosli, G., Canu, S., Ong, C., (2015). For the PCVM the provided kernel evaluations are considered only as basis functions and no mercer conditions are implied. In contrast to the iKFD the PCVM is a *sparse* probabilistic kernel classifier pruning unused basis functions during training, applicable to arbitrary positive definite *and* indefinite kernel matrices. A recent review about learning with indefinite proximities can be found in Schleif and Tino (2015).

While being very efficient these methods do not scale to larger datasets with in general cubic complexity. In Schleif et al. (2015a); Gisbrecht and Schleif (2015) the authors proposed a few Nyström based (see e.g. Williams and Seeger (2000)) approximation techniques to improve the scalability of the PCVM for low rank matrices. The suggested techniques use the Nyström approximation in a non-trivial way to provide *exact* eigenvalue estimations also for *indefinite* kernel matrices. This approach is very generic and can be applied in different algorithms. In this contribution we further extend our previous work and not only derive a low rank approximation of the indefinite kernel Fisher discriminant, but also address the landmark selection from a novel view point. The obtained Ny-iKFD approach is linear in runtime and memory consumption for low rank matrices. The formulation is exact if the rank of the matrix equals the number of independent landmarks points. The selection of the landmarks of the Nyström approximation is a critical point addressed in previous work (see e.g. Zhang and Kwok (2010); Si et al. (2014); Brabanter et al. (2010)). In general these strategies use the full psd kernel matrix or expect that the kernel is of some standard class like an RBF kernel. In each case the approaches presented so far are costly in runtime and memory consumption as can be seen in the subsequent experiments.

Additionally, former approaches for landmark selection aim on generic matrix reconstructions of positive semi definite (psd) kernels. We propose a restricted reconstruction of the psd or non-psd kernel matrix with respect to a *supervised* learning scenario only. We do not any longer expect to obtain an accurate kernel

reconstruction from the approximated matrix (e.g. by using the Frobenius norm) but are pleased if the approximated matrix preserves the class boundaries in the data space.

In Gisbrecht and Schleif (2015) the authors derived methods to approximate large proximity matrices by means of the Nyström approximation and conversion rules between similarities and dissimilarities. These techniques have been applied in Schleif et al. (2015a) and Schleif, F.-M., Gisbrecht, A., Tino, P., (2015b) in a proof of concept setting, to obtain approximate models for the Probabilistic Classification Vector Machine and the Indefinite Fisher Kernel Discriminant analysis using a random landmark selection scheme. This work is substantially extended and detailed in this article with a specific focus on *indefinite* kernels, only. A novel landmark selection scheme is proposed. Based on this new landmark selection scheme we provide detailed new experimental results and compare to alternative landmark selection approaches.

Structure of the paper: First we give some basic notations necessary in the subsequent derivations. Then we review iKFD and PCVM as well as some approximation concepts proposed by the authors in Schleif et al. (2015a) which are based on the well known Nyström approximation. Subsequently, we consider the landmark selection problem in more detail and show empirically results motivating a supervised selection strategy. Finally we detail the reformulation of iKFD and PCVM based on the introduced concepts and show the efficiency in comparison to Ny-PCVM and Ny-iKFD for various indefinite proximity benchmark data sets.

2. Methods

2.1. Notation and basic concepts

Let $\mathcal{X}$ be a collection of N objects x_i, $i = 1, 2, ..., N$, in some input space. If the similarity function or inner product used to compare two objects x_i, x_j is metric, proper mercer kernels can be obtained as discussed subsequently. A classical similarity function in this context, is the Euclidean inner product with the respective Euclidean distance which is a frequent core component of various metric kernel functions, like the famous radial basis function (rbf) kernel.

Now, let $\phi : \mathcal{X} \mapsto \mathcal{H}$ be a mapping of patterns from $\mathcal{X}$ to a high-dimensional or infinite dimensional Hilbert space $\mathcal{H}$ equipped with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$. The transformation ϕ is in general a non-linear mapping to a high-dimensional space $\mathcal{H}$ and may in general not be given in an explicit form. Instead a kernel function $k : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$ is given which encodes the inner product in $\mathcal{H}$. The

kernel k is a positive (semi) definite function such that $k(x, x') = \phi(x)^\top \phi(x')$ for any $x, x' \in \mathcal{X}$. The matrix $K := \Phi^\top \Phi$ is an $N \times N$ kernel matrix derived from the training data, where $\Phi : [\phi(x_1), \ldots, \phi(x_N)]$ is a matrix of images (column vectors) of the training data in $\mathcal{H}$. The motivation for such an embedding comes with the hope that the non-linear transformation of input data into higher dimensional $\mathcal{H}$ allows for using linear techniques in $\mathcal{H}$. Kernelized methods process the embedded data points in a feature space utilizing only the inner products $\langle \cdot, \cdot \rangle_\mathcal{H}$ (kernel trick) (Shawe-Taylor and Cristianini, 2004), without the need to explicitly calculate ϕ. The specific kernel function can be very generic. Most prominent are the linear kernel with $k(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$ where $\langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$ is the Euclidean inner product or the rbf kernel $k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{||\mathbf{x}-\mathbf{x}'||^2}{2\sigma^2}\right)$, with σ as a free parameter. Thereby it is assumed that the kernel function $k(\mathbf{x}, \mathbf{x}')$ is positive semi definite (psd). This assumption is however not always fulfilled, and the underlying similarity measure may not be metric and hence not lead to a mercer kernel. Examples can be easily found in domain specific similarity measures as mentioned before and detailed later on. These measures imply *indefinite* kernels. In what follows we will review some basic concepts and approaches related to such non-metric situations.

2.2. Krein and Pseudo-Euclidean spaces

A Krein space is an *indefinite* inner product space endowed with a Hilbertian topology.

Definition 1 (Inner products and inner product space). *Let $\mathcal{K}$ be a real vector space. An inner product space with an indefinite inner product $\langle \cdot, \cdot \rangle_\mathcal{K}$ on $\mathcal{K}$ is a bi-linear form where all $f, g, h \in \mathcal{K}$ and $\alpha \in \mathbb{R}$ obey the following conditions.*

- *Symmetry:* $\langle f, g \rangle_\mathcal{K} = \langle g, f \rangle_\mathcal{K}$
- *linearity:* $\langle \alpha f + g, h \rangle_\mathcal{K} = \alpha \langle f, h \rangle_\mathcal{K} + \langle g, h \rangle_\mathcal{K}$*;*
- $\langle f, g \rangle_\mathcal{K} = 0 \; \forall g \in K$ *implies* $f = 0$

An inner product is positive definite if $\forall f \in \mathcal{K}$, $\langle f, f \rangle_\mathcal{K} \geq 0$, negative definite if $\forall f \in \mathcal{K}$, $\langle f, f \rangle_\mathcal{K} \leq 0$, otherwise it is indefinite. A vector space $\mathcal{K}$ with inner product $\langle \cdot, \cdot \rangle_\mathcal{K}$ is called an inner product space.

Definition 2 (Krein space and pseudo Euclidean space). *An inner product space* $(\mathcal{K}, \langle \cdot, \cdot \rangle_{\mathcal{K}})$ *is a Krein space if we have two Hilbert spaces* $\mathcal{H}_+$ *and* $\mathcal{H}_-$ *spanning* $\mathcal{K}$ *such that* $\forall f \in \mathcal{K}$ *we have* $f = f_+ + f_-$ *with* $f_+ \in \mathcal{H}_+$ *and* $f_- \in \mathcal{H}_-$ *and* $\forall f, g \in \mathcal{K}$, $\langle f, g \rangle_{\mathcal{K}} = \langle f_+, g_+ \rangle_{\mathcal{H}_+} - \langle f_-, g_- \rangle_{\mathcal{H}_-}$. *A finite-dimensional Krein-space is a so called pseudo Euclidean space (pE).*

Indefinite kernels are typically observed by means of domain specific non-metric similarity functions (such as alignment functions used in biology (Smith et al., 1981)), by specific kernel functions - e.g. the Manhattan kernel $k(\mathbf{x}, \mathbf{x}') = -||\mathbf{x} - \mathbf{x}'||_1$, tangent distance kernel (Haasdonk and Keysers, 2002) or divergence measures plugged into standard kernel functions (Cichocki and Amari, 2010). Another source of non-psd kernels are noise artifacts on standard kernel functions (Haasdonk, 2005).

For such spaces vectors can have negative squared "norm", negative squared "distances" and the concept of orthogonality is different from the usual Euclidean case. In the subsequent experiments our input data are in general given by a symmetric indefinite kernel matrix K.

Given a symmetric *dissimilarity* matrix with zero diagonal [1], an embedding of the data in a pseudo-Euclidean vector space determined by the eigenvector decomposition of the associated similarity matrix $\mathbf{S}$ is always possible (Goldfarb, 1984) [2]

Given the eigendecomposition of $\mathbf{S}$, $\mathbf{S} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$, we can compute the corresponding vectorial representation $\mathbf{V}$ in the pseudo-Euclidean space by

$$\mathbf{V} = \mathbf{U}_{p+q+z} \left| \mathbf{\Lambda}_{p+q+z} \right|^{1/2} \tag{1}$$

where $\mathbf{\Lambda}_{p+q+z}$ consists of p positive, q negative non-zero eigenvalues and z zero eigenvalues. $\mathbf{U}_{p+q+z}$ consists of the corresponding eigenvectors. The triplet (p, q, z) is also referred to as the signature of the Pseudo-Euclidean space. This operation is however very costly and should be avoided for larger data sets. A detailed presentation of similarity and dissimilarity measures, and mathematical aspects of metric and non-metric spaces is provided in (Pekalska and Duin, 2005).

[1] A similarity matrix can be easily converted into squared dissimilarities using $d^2(x, y) = k(x, x) + k(y, y) - 2 \cdot k(x, y)$.

[2] The associated similarity matrix can be obtained by double centering (Pekalska and Duin, 2005) of the (squared) dissimilarity matrix. $\mathbf{S} = -\mathbf{J}\mathbf{D}\mathbf{J}/2$ with $\mathbf{J} = (\mathbf{I} - \mathbf{1}\mathbf{1}^\top/N)$ and identity matrix $\mathbf{I}$ and vector of ones $\mathbf{1}$.

2.3. Indefinite Fisher and kernel quadratic discriminant

In Haasdonk and Pekalska (2008); Pekalska and Haasdonk (2009) the indefinite kernel Fisher discriminant analysis (iKFD) and indefinite kernel quadratic discriminant analysis (iKQD) was proposed focusing on binary classification problems, recently extended by a weighting scheme in Yang and Fan (2013)[3].

The initial idea is to embed the training data into a Krein space (see Def. 2) and to apply a modified kernel Fisher discriminant analysis or kernel quadratic discriminant analysis for indefinite kernels. Consider binary classification and a data set of input-target training pairs $D = \{\mathbf{x}_i, y_i\}_{i=1}^N$, where $y_i \in \{-1, +1\}$. Given the indefinite kernel matrix K and the embedded data in a pseudo-Euclidean space (pE), the linear Fisher Discriminant function $f(x) = \langle w, \Phi(\mathbf{x})\rangle_{pE} + b$ is based on a weight vector $\mathbf{w}$ such that the between-class scatter is maximized while the within-class scatter is minimized along w. $\Phi(\mathbf{x})$ is a vector of basis function evaluations for data item $\mathbf{x}$ and b is a bias term. This direction is obtained by maximizing the Fisher criterion in the pseudo Euclidean space

$$J(\mathbf{w}) = \frac{\langle \mathbf{w}, \Sigma_{pE}^b \mathbf{w}\rangle_{pE}}{\langle \mathbf{w}, \Sigma_{pE}^w \mathbf{w}\rangle_{pE}}$$

where $\Sigma_{pE}^b = \Sigma_b J$ is the scatter matrix in the pseudo Euclidean space, with $J = diag(\mathbf{1}_p, -\mathbf{1}_q)$ where $\mathbf{1}_p \in \mathbb{R}^p$ denotes the p-dimensional vector of all ones. The number of positive eigenvalues is denoted by p and for the negative eigenvalues by q. The *within-scatter-matrix* in the pseudo-Euclidean space is given as $\Sigma_{pE}^w = \Sigma_w J$. The *Euclidean* between- and within-scatter-matrices can be expressed as:

$$\Sigma_b = (\mu_+ - \mu_-)(\mu_+ - \mu_-)^\top \quad (2)$$

$$\Sigma_w = \sum_{i \in I_+} (\phi(\mathbf{x}_i) - \mu_+)(\phi(\mathbf{x}_i) - \mu_+)^\top + \sum_{i \in I_-} (\phi(\mathbf{x}_i) - \mu_-)(\phi(\mathbf{x}_i) - \mu_-)^\top \quad (3)$$

Where the set of indices of each class are $I_+ := \{i : y_i = +1\}$ and $I_- := \{i : y_i = -1\}$. In Haasdonk and Pekalska (2008) it is shown that the Fisher Discriminant in the pE space $\in \mathbb{R}^{(p,q)}$ is identical to the Fisher Discriminant in the associated Euclidean space $\mathbb{R}^{p+q}$. To avoid the explicit embedding into the pE space a kernelization is considered such that the weight vector $w \in \mathbb{R}^{(p,q)}$ is expressed as a linear combination of the training data $\phi(\mathbf{x}_i$, hence $w = \sum_{i=1}^N \alpha_i \phi(\mathbf{x}_i$. Transferred to the Fisher criterion this allows to use the kernel trick. A similar strategy

[3]For multiclass problems a classical 1 vs rest wrapper is used within this paper

can be used for KQD as well as the indefinite kernel PCA Pekalska and Haasdonk (2009).

2.4. Probabilistic Classification Vector Learning

PCVM uses a kernel regression model $\sum_{i=1}^{N} w_i \phi_i(\mathbf{x}) + b$ with a link function, with w_i being again the weights of the basis functions $\phi_i(\mathbf{x})$ and b as a bias term. The Expectation Maximization (EM) implementation of PCVM Chen et al. (2014) uses the probit link function, i.e. $\Psi(x) = \int_{-\infty}^{x} \mathcal{N}(t|0,1)dt$,where $\Psi(x)$ is the cumulative distribution of the normal distribution $\mathcal{N}(0,1)$. We get: $l(\mathbf{x};\mathbf{w},b) = \Psi\left(\sum_{i=1}^{N} w_i\phi_i(\mathbf{x}) + b\right) = \Psi\left(\Phi(\mathbf{x})\mathbf{w} + b\right)$

In the PCVM formulation Chen et al. (2009a), a truncated Gaussian prior N_t with support on $[0,\infty)$ and mode at 0 is introduced for each weight w_i and a zero-mean Gaussian prior is adopted for the bias b. The priors are assumed to be mutually independent. $p(\mathbf{w}|\alpha) = \prod_{i=1}^{N} p(w_i|\alpha_i) = \prod_{i=1}^{N} N_t(w_i|0,\alpha_i^{-1})$, $p(b|\beta) = \mathcal{N}(b|0,\beta^{-1}), \delta(x) = \mathbf{1},\ x > 0$.

$$p(w_i|\alpha_i) = \begin{cases} 2\mathcal{N}(w_i|0,\alpha_i^{-1}) & \text{if } y_i w_i > 0 \\ 0 & \text{otherwise} \end{cases} = 2\mathcal{N}(w_i|0,\alpha_i^{-1}) \cdot \delta(y_i w_i).$$

We follow the standard probabilistic formulation and assume that $z(\mathbf{x}) = \Phi(\mathbf{x})\mathbf{w} + b$ is corrupted by an additive random noise ϵ , where $\epsilon \sim \mathcal{N}(0,1)$. According to the probit link model, we have:

$$h(\mathbf{x}) = \Phi(\mathbf{x})\mathbf{w} + b + \epsilon \geq 0, y = 1, \quad h(\mathbf{x}) = \Phi(\mathbf{x})\mathbf{w} + b + \epsilon < 0, y = -1 \quad (4)$$

and obtain: $p(y = 1|\mathbf{x},\mathbf{w},b) = p(\Phi(\mathbf{x})\mathbf{w} + b + \epsilon \geq 0) = \Psi(\Phi(\mathbf{x})\mathbf{w} + b)$. $h(\mathbf{x})$ is a latent variable because ϵ is an unobservable variable. We collect evaluations of $h(\mathbf{x})$ at training points in a vector $\mathbf{H}(\mathbf{x}) = (h(\mathbf{x_1}), \ldots, h(\mathbf{x_N}))^\top$. In the expectation step the expected value $\bar{\mathbf{H}}$ of $\mathbf{H}$ with respect to the posterior distribution over the latent variables is calculated (given old values $\mathbf{w}^{\text{old}}, b^{\text{old}}$). In the maximization step the parameters are updated through

$$\mathbf{w}^{\text{new}} = M(M\Phi^\top(\mathbf{x})\Phi(\mathbf{x})M + I_N)^{-1} M(\Phi^\top(\mathbf{x})\bar{\mathbf{H}} - b\Phi^\top(\mathbf{x})\mathbf{I}) \quad (5)$$

$$\mathbf{b}^{\text{new}} = t(1 + tNt)^{-1} t(\mathbf{I}^\top\bar{\mathbf{H}} - \mathbf{I}^\top\Phi(\mathbf{x})\mathbf{w}) \quad (6)$$

where I_N is a N-dimensional identity matrix and $\mathbf{I}$ a all-ones vector, the diagonal elements in the diagonal matrix M are:

$$m_i = (\bar{\alpha}_i)^{-1/2} = \begin{cases} \sqrt{2}w_i & \text{if} \quad y_i w_i \geq 0 \\ 0 & \text{else} \end{cases} \quad (7)$$

and the scalar $t = \sqrt{2}|b|$. For further details can be found in Chen et al. (2009a). Even though kernel machines and their derivatives have shown great promise in practical application, their scope is somehow limited by the fact that the computational complexity grows rapidly with the size of the kernel matrix (number of data items). Among methods suggested to deal with this issue in the literature, the Nyström method has been popular and widely used.

3. Nyström approximated matrix processing

The Nyström approximation technique has been proposed in the context of kernel methods in (Williams and Seeger, 2000). Here, we give a short review of this technique before it is employed in PCVM and iKFD. One well known way to approximate a $N \times N$ Gram matrix, is to use a low-rank approximation. This can be done by computing the eigendecomposition of the kernel matrix $K = U\Lambda U^T$, where U is a matrix, whose columns are orthonormal eigenvectors, and Λ is a diagonal matrix consisting of eigenvalues $\Lambda_{11} \geq \Lambda_{22} \geq ... \geq 0$, and keeping only the m eigenspaces which correspond to the m largest eigenvalues of the matrix. The approximation is $\tilde{K} \approx U_{N,m}\Lambda_{m,m}U_{m,N}$, where the indices refer to the size of the corresponding submatrix restricted to the larges m eigenvalues. The Nyström method approximates a kernel in a similar way, without computing the eigendecomposition of the whole matrix, which is an $O(N^3)$ operation.

By the Mercer theorem kernels $k(\mathbf{x}, \mathbf{x}')$ can be expanded by orthonormal eigenfunctions φ_i and non negative eigenvalues λ_i in the form

$$k(\mathbf{x}, \mathbf{x}') = \sum_{i=1}^{\infty} \lambda_i \varphi_i(\mathbf{x}) \varphi_i(\mathbf{x}').$$

The eigenfunctions and eigenvalues of a kernel are defined as the solution of the integral equation

$$\int k(\mathbf{x}', \mathbf{x}) \varphi_i(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} = \lambda_i \varphi_i(\mathbf{x}'),$$

where $p(\mathbf{x})$ is the probability density of $\mathbf{x}$. This integral can be approximated based on the Nyström technique by an i.i.d. sample $\{\mathbf{x}^k\}_{k=1}^m$ from $p(\mathbf{x})$:

$$\frac{1}{m} \sum_{k=1}^{m} k(\mathbf{x}', \mathbf{x}^k) \varphi_i(\mathbf{x}^k) \approx \lambda_i \varphi_i(\mathbf{x}').$$

Using this approximation we denote with $K^{(m)}$ the corresponding $m \times m$ Gram sub-matrix and get the corresponding matrix eigenproblem equation as:

$$K^{(m)}U^{(m)} = U^{(m)}\Lambda^{(m)}$$

with $U^{(m)} \in \mathbb{R}^{m\times m}$ is column orthonormal and $\Lambda^{(m)}$ is a diagonal matrix.

Now we can derive the approximations for the eigenfunctions and eigenvalues of the kernel k

$$\lambda_i \approx \frac{\lambda_i^{(m)} \cdot N}{m}, \quad \varphi_i(\mathbf{x}') \approx \frac{\sqrt{m/N}}{\lambda_i^{(m)}} \mathbf{k}_x'^{,\top} \mathbf{u}_i^{(m)}, \tag{8}$$

where $\mathbf{u}_i^{(m)}$ is the ith column of $U^{(m)}$. Thus, we can approximate φ_i at an arbitrary point $\mathbf{x}'$ as long as we know the vector $\mathbf{k}'_x = (k(\mathbf{x}^1, \mathbf{x}'), ..., k(\mathbf{x}^m, \mathbf{x}'))$. For a given $N \times N$ Gram matrix K one may randomly choose m rows and respective columns. The corresponding indices are called landmarks, and should be chosen such that the, data distribution is sufficiently covered. Strategies how to chose the landmarks have recently been addressed in Zhang and Kwok (2010); Zhang et al. (2008) and Gittens and Mahoney (2013); Brabanter et al. (2010). We denote these rows by $K_{m,N}$. Using the formulas Eq. (8) we obtain $\tilde{K} = \sum_{i=1}^{m} 1/\lambda_i^{(m)} \cdot K_{m,N}^T (\mathbf{u}_i^{(m)})^T (\mathbf{u}_i^{(m)}) K_{m,N}$, where $\lambda_i^{(m)}$ and $\mathbf{u}_i^{(m)}$ correspond to the $m \times m$ eigenproblem. Thus we get, $K_{m,m}^{-1}$ denoting the Moore-Penrose pseudoinverse,

$$\tilde{K} = K_{N,m} K_{m,m}^{-1} K_{m,N}. \tag{9}$$

as an approximation of K. This approximation is exact, if $K_{m,m}$ has the same rank as K.

3.1. Pseudo Inverse and Singular Value Decomposition of a Nyström approximated matrix

In the Ny-PCVM approach discussed in Section 5 we need the pseudo inverse of a Nyström approximated matrix while for the Ny-iKFD a Nyström approximated eigenvalue decomposition (EVD) is needed.

A Nyström approximated pseudo inverse can be calculated by a modified singular value decomposition (SVD) with a rank limited by $r^* = \min\{r, m\}$ where r is the rank of the pseudo inverse and m the number of landmark points. The output is given by the rank reduced left and right singular vectors and the reciprocal of the singular values. The singular value decomposition based on a Nyström approximated similarity matrix $\tilde{K} = K_{Nm} K_{m,m}^{-1} K_{Nm}^{\top}$ with m landmarks, calculates

the left singular vectors of $\tilde{K}$ as the eigenvectors of $\tilde{K}\tilde{K}^\top$ and the right singular vectors of $\tilde{K}$ as the eigenvectors of $\tilde{K}^\top\tilde{K}$[4]. The non-zero singular values of $\tilde{K}$ are then found as the square roots of the non-zero eigenvalues of both $\tilde{K}^\top\tilde{K}$ or $\tilde{K}\tilde{K}^\top$. Accordingly one only has to calculate a new Nyström approximation of the matrix $\tilde{K}\tilde{K}^\top$ using e.g. the same landmark points as for the input matrix $\tilde{K}$. Subsequently an eigenvalue decomposition (EVD) is calculated on the approximated matrix $\zeta = \tilde{K}\tilde{K}^\top$. For a matrix approximated by Eq. (9) it is possible to compute its exact eigenvalue estimators in linear time.

3.2. Eigenvalue decomposition of a Nyström approximated matrix

To compute the eigenvectors and eigenvalues of an *indefinite* matrix we first compute the squared form of the Nyström approximated kernel matrix. Let K be a psd similarity matrix, for which we can write its decomposition as

$$\tilde{K} = K_{N,m}K_{m,m}^{-1}K_{m,N} = K_{N,m}U\Lambda^{-1}U^\top K_{N,m}^\top \qquad = BB^\top,$$

where we defined $B = K_{N,m}U\Lambda^{-1/2}$ with U and Λ being the eigenvectors and eigenvalues of $K_{m,m}$, respectively.
Further it follows for the *squared* $\tilde{K}$:

$$\tilde{K}^2 = BB^\top BB^\top = BVAV^\top B^\top,$$

where V and A are the eigenvectors and eigenvalues of $B^\top B$, respectively. The square operation does not change the eigenvectors of K but only the eigenvalues. The corresponding eigenequation can be written as $B^\top Bv = av$. Multiplying with B from left we get: $\underbrace{BB^\top}_{\tilde{K}}\underbrace{(Bv)}_{u} = a\underbrace{(Bv)}_{u}$. It is clear that A must be the matrix with the eigenvalues of $\tilde{K}$. The matrix Bv is the matrix of the corresponding eigenvectors, which are orthogonal but not necessary orthonormal. The normalization can be computed from the decomposition:

$$\tilde{K} = B\underbrace{VV^\top}_{diag(1)}B^\top = BVA^{-1/2}AA^{-1/2}V^\top B^\top = CAC^\top,$$

where we defined $C = BVA^{-1/2}$ as the matrix of orthonormal eigenvectors of K. The eigenvalues of $\tilde{K}$ can be obtained using $A = C^\top\tilde{K}C$. Using this derivation

[4]For symmetric matrices we have $\tilde{K}\tilde{K}^\top = \tilde{K}^\top\tilde{K}$

we can obtain exact eigenvalues and eigenvectors of an indefinite low rank kernel matrix K, given rank$(K) = m$ and the landmarks points are independent[5]

The former approximation scheme is focused on preserving the full low rank eigen structure of the underlying data space. The accuracy of this approximation is typically measured by the Frobenius norm. A low value of the Frobenius norm of the approximated versus the original kernel matrix ensures that the approximated kernel matrix $\tilde{K}$ can be used similar as K for any kernel based data analysis method like kernel-PCA, kernel-k-means, SVM, laplacian eigenmaps, preserving also small between point distances. In the context of classification this requirement is very strong and unnecessary. We suggest to restrict the approximation $\tilde{K}$ to a low rank kernel which preserves *only* the between class distances focusing on class separation. To achieve this objective we suggest to use a supervised landmark selection scheme. This is introduced in the following section and compared with a number of baseline methods.

4. Supervised landmark selection using minimum enclosing balls

The Nyström approximation is based on m characteristic landmark points taken from the dataset. The number of landmarks should be sufficiently large and the landmarks should be diverse enough to get accurate approximations of the dominating singular vectors of the similarity matrix. In Zhang and Kwok (2010) multiple strategies for landmark selection have been studied and a clustering based approach was suggested to find the specific landmarks. Thereby the number of landmarks is a user defined parameter and a classical kmeans algorithm is applied on the kernel matrix to identify characteristic landmark points in the empirical feature space. This approach is quite effective (see Zhang and Kwok (2010)), with some small improvements using an advanced clustering scheme as shown in Si et al. (2014). We will use it as a baseline for an advanced landmark section approach. Further we will also consider a pure random selection strategy. It should be noted that the formulation given in Zhang and Kwok (2010) takes the full kernel matrix as an input into the k-means clustering. This is obviously also very costly and may become inapplicable for larger kernel matrices [6] It is however not yet

[5] An implementation of this linear time eigen-decomposition for low rank indefinite matrices is available at: `http://www.techfak.uni-bielefeld.de/~fschleif/eigenvalue_corrections_demos.tgz`.

[6] It may however be possible to circumvent this full complexity approach e.g. by subsampling concepts or by more advanced concepts of k-means, but this is not the focus of this paper.

clear how the *number* of landmarks can be appropriately chosen. If the number of landmarks is large we can expect the data space to be sufficiently covered after the clustering but the model complexity can become prohibitive. On the other hand if the number of landmarks is to small the clustering may lead to inappropriate results by merging disjunct parts of the data space. We propose to consider the Nyström approximation in a restricted form with respect to a *supervised* learning problem only. This relieves us from the need of a perfect reconstruction of the kernel matrix. It is in fact sufficient to reconstruct the kernel such that it is close to the ideal kernel (see e.g. Kwok and Tsang (2003)). We will however not learn an idealized kernel as proposed in Kwok and Tsang (2003), which by itself is very costly for large scale matrices, but provide a landmark selection strategy leading into a similar direction.

Typically the approximation quality of a Nyström approximated similarity matrix is evaluated using the Frobenius norm.For real valued data the Frobenius norm of two squared matrices, is simply the sum of the squared difference between the individual kernel entries. The Frobenius norm is very sensitive to small perturbations of the kernel matrix, addressing also small local geometric violations. Instead we propose a margin based similarity measure between matrices taking labels into account.

We define a supervised similarity measure between two matrices to estimate improved prediction performance if a *linear* classifier is used [7]

Definition 3 (Supervised matrix similarity score (SMSS)). *Assume we have a two class problem with* $y \in \{-1, 1\}$[8]. *Given a similarity function* $K : X \times X \to \mathbb{R}$ *for a learning problem* P *with underlying points* $(x, y) \sim P$. *We define a score* $s(\hat{K}, K) = \frac{f(\hat{K})}{f(K)}$ *with*
$f(S) = \sum_{y \in \{-1,1\}} \left| E_{(x_i, x'_j)|(y_i = y)}\{S(x_i, x'_j)|y = y'_j\} - E_{(x_i, x'_j)|y_i = y}\{S(x_i, x'_j)|y \neq y'_j\} \right|$

The function $f(S)$ provides a margin estimate of the scores between pairwise similarities of a class y and pairwise similarities between the entries of class y with respect to entries of the other class(es). Assuming that $f(\hat{K})$ gives a margin estimate of the improved matrix and $f(K)$ for the original input matrix, the score

[7] Note that for non-linear separable data the kernel trick can be used and we can still use linear decision functions as shown in iKFD and PCVM.

[8] The extension to more than two classes is straight forward

$s(\hat{K}, K) > 1$, if the margin has increased and $s(\hat{K}, K) \leq 1$ otherwise. Similar as in Balcan et al. (2008) we assume that a good linear classifier can be obtained from a given similarity function (or kernel) if the similarities between classes are much lower than those within the same class. However the score $s(\hat{K}, K)$ is only a rough estimate. It will likely work well for datasets which can easily be modeled by conceptually related classifiers focusing e.g. on exemplar based representations. The median classifier proposed in Nebel et al. (2015) is such a simple classifier. In its simplest form it identifies for each class a single basis function or prototype, showing maximum margin with respect to the other prototypes with different class labels. In Nebel et al. (2015) it was shown that such a classifier can be very efficient also for a variety of classification problems.

In our study the modified similarity matrix $\hat{K}$ is a Nyström approximated matrix, where the landmarks should be chosen to keep good prediction accuracy instead of a good data reconstruction, as typically aimed for. Note that due to the approximation it may in fact happen that the SMSS values is below 1, indicating a decreased discrimination power with respect to the full original matrix. But in the considered setting the approximation is a mandatory step and we try to achieve a large SMSS value for the approximated similarity matrix.

Apparently the (supervised) representation accuracy of the Nyström approximation of K depends on the number and type of the selected landmarks. We propose to calculate minimum enclosing ball solutions (MEB) on the individual classwise kernels, to address both problems:

1. finding a sufficient number of landmarks
2. find landmarks explaining the data characteristics and preserve a good class separation for $\hat{K}$

As an additional constraint we are looking for an approach where also indefinite proximity matrices can be processed without costly preprocessing steps.

4.1. MEB for psd input kernels

We denote the set of indices or points of a sub kernel matrix referring to class j by R_j. Assuming approximately spherical clusters, we can approximate this problem by the **minimum enclosing ball** :

$$\begin{array}{ll} \min_{R^2, w_j} & R^2 \\ \text{such that} & \|w_j - \Phi(\xi_i)\|^2 \leq R^2 \quad \forall \xi_i \in R_j \end{array}$$

where R is the radius of the sphere and w_j is a center point which can be indirectly represented in the kernel space) as a weighted linear combination of the points in

R_j. The assumption of a sphere is in fact no substantial restriction if the provided kernel is sufficiently expressive. This is also the reason why core-vector data description (CVDD) can be used as a linear time replacement for support vector data description Tsang, I. W., Kwok, J. T., Cheung, P., (2005).

It has been shown e.g. in Badoiu and Clarkson (2008) that the minimum enclosing ball can be approximated with quality ϵ in (worst case) linear time using an algorithm which requires only a constant subset of the receptive field R_j, the core set. Given fixed quality ϵ, the following algorithm converges in $\mathcal{O}(1/\epsilon^2)$ steps:

MEB:
$S := \{\xi_i, \xi_k\}$ for a pair of largest distance $\|\Phi(\xi_i) - \Phi(\xi_k)\|^2$ in R_j and ξ_i chosen randomly
repeat
  solve MEB$(S) \rightarrow \tilde{w}_j, R$
  if exists $\xi_l \in R_j$ where $\|\Phi(\xi_l) - \tilde{w}_j\|^2 > R^2(1+\epsilon)^2$ **then**
    $S := S \cup \{\xi_l\}$
  end if
until all ξ_l are covered by the $R(1+\epsilon)$ ball in the feature space
return $\tilde{w}_j$

In each step, the MEB problem is solved for a small subset of constant size only. This is possible by referring to the dual problem which has the form

$$\begin{array}{ll} \min_{\alpha_i \geq 0} & \sum_{ij} \alpha_i \alpha_j k_{ij} - \sum_i \alpha_i k_{ii}^2 \\ \text{where} & \sum_i \alpha_i = 1 \end{array}$$

with data points occurring only as dot products, i.e. kernelization is possible. The same holds for all distance computations of the approximate MEB problem. Note that the dual MEB problem provides a solution in terms of the dual variables α_i. The identified finite number of core points (those with non-vanishing α_i) will be used as landmarks for this class and considered to be sufficient to represent the enclosing sphere of the data. Each class is represented by at least two core points. Combining all core sets of the various classes provides us with the full set of landmarks used to get a Nyström approximation of K.

The MEB solution typically consists of a very small number of points (independent of N), sufficient to describe the hyper-ball enclosing the respective data. If the kernel is psd we can use the MEB approach directly in the kernel space.

4.2. MEB for non-psd input kernels

If the given kernel is non-psd we either can apply various eigenvalue correction approaches see Schleif and Tino (2015) or we use $\hat{K} = K \cdot K^\top$, which can

also be easily done for Nyström approximated matrices without calculating a full matrix (see first part of Eq. (15)). This procedure does not change the eigenvectors of K but takes the square of the eigenvalues such that $\hat{K}$ becomes psd. It should be noted that if we use $\hat{K}$ as an input of a kernel k-means algorithm this is equivalent as using K as the input of the classical k-means with Euclidean distance as suggested in Zhang and Kwok (2010).

The proposed supervised landmark selection using MEB does not only identify a good estimate for the number of landmarks but also ensures that the landmarks are sufficient to explain the data space. The solutions of the MEB consist of non-redundant points at the perimeter of the sphere, which can considered to be unrelated, although not necessarily orthogonal in the similarity space (with potentially squared negative eigenvalues). Especially only those points are included in the MEB solution which are needed to explain the sphere such that redundancy within this set is avoided Badoiu and Clarkson (2008). Therefore for each class the MEB solutions provides a local span of the underlying eigen-space. The combination of the different subspaces can lead to redundancy but we can expect that the full data space is sufficiently covered. We will show the effectiveness of this approach in some short experiments.

4.3. Small scale experiments - landmark selection scheme

We use the ball dataset as proposed in Duin and Pekalska (2010). It is an artificial dataset based on the surface distances of randomly positioned balls of two classes having a slightly different radius. The dataset is non-Euclidean with substantial information encoded in the negative part of the eigenspectrum. We generated the data with 100 samples per class leading to an $N \times N$ dissimilarity matrix $\mathbf{D}$, with $N = 200$.

We also use the protein data (213 pts, 4 classes) set represented by an indefinite similarity matrix, with a high intrinsic dimension Schleif and Tino (2015). Further we analyzed two simulated metric datasets which are not linear separable using the Euclidean norm: (1) the checker board data, generated as a two dimensional dataset with datapoints organized on a 3×3 checkerboard, with alternating labels. This dataset has *multi-modal* classes. (2) a simple gaussian cloud dataset with two gaussian with substantial overlap. The simulated data have been represented by an extreme learning machine (elm) kernel. Checker is linear separable in the elm-kernel space, whereas Gaussian is not separable by construction.

It should be noted that the elm kernel, used for the vectorial data, typically increases the number of non-vanishing eigenvalues such that the original two dimensional data are finally indeed higher dimensional and not representable by

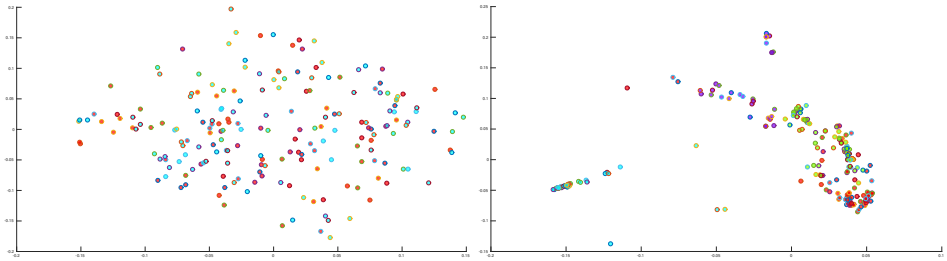

(a) Ball dataset

(b) Protein dataset

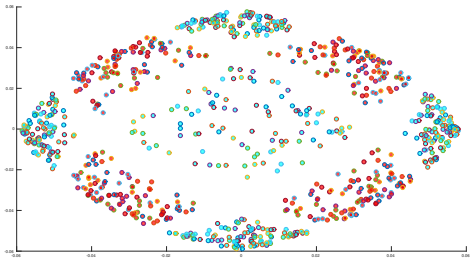

(c) Checker dataset

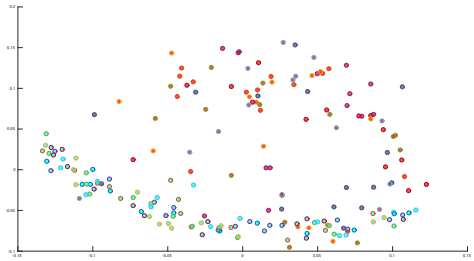

(d) Gauss dataset

Figure 1: Laplacian eigenmap visualization of the initial test and simulated similarity matrices using $K \cdot K^\top$. Colors/shades indicate the different classes. Axis labeling is arbitrary.

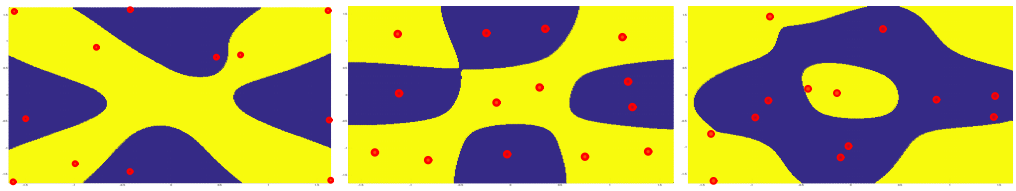

(a) Checker board data with the MEB selection scheme.

(b) Checker board data with the k-means selection scheme using #MEB landmarks.

(c) Checker board data with the random selection scheme using #MEB landmarks.

Figure 2: Typical plots of the checker board data - taken from the crossvalidation models - with iKFD predictions using different landmark selection schemes and an elm kernel. The worst result $\approx 72\%$ is obtained by plot c) using the random sampling strategy whereby the number of landmarks was chosen from the MEB approach. The selected landmark points are indicated as (red) circles. In plot b) one clearly sees that k-means has rearranged the points to cover the whole data space. For the random approach we observe that some points are very close to each other (and have the same label) and are therefore not very informative.

only two basis functions. Two dimensional visualizations of the unapproximated $K \cdot K^\top$ similarity matrices obtained by using laplacian eigenmaps Belkin and Niyogi (2003). are shown in Figure 1.For the checker board data we also show two-dimensional plots of the obtained iKFD decision boundaries and different landmark selection schemes in Figure 2.

Now the obtained (indefinite) kernel matrix has been used in the iKFD in six different ways using different landmark selection schemes:

a) we used the original kernel matrix (SIM1),

b) the matrix is Nystöm approximated using the MEB approach (SIM2),

c) the matrix is Nyström approximated using the approach of Zhang and Kwok (2010) where the number of landmarks is taken from the MEB solution (SIM3),

d) using the approach of Zhang and Kwok (2010) but with C landmarks where C is the number of classes (SIM4)

e) using a random sample of C landmarks (SIM5). SIM5 can be considered as a very basic baseline approach.

f) using an entropy based selection as proposed in Brabanter et al. (2010) (SIM6) [9] where the number of landmarks is again taken from the MEB solution

One may also simply use a *very* large number of randomly selected landmarks, but this can become prohibitive if N is large such that the calculation of $N \times m$ similarities can be costly in memory and runtime. Further it can be very unattractive to have a larger m for the out of sample extension to new points. If for example costly alignment scores are used one is interested on having a very small m to avoid large costs in the test phase of the model.

The results of a 10-fold crossvalidation are shown in the Table 1 with runtimes given in Table 2. Here and in the following experiments the landmark selection was part of the crossvalidation scheme and the landmarks are selected on the training set only and the test data have been mapped to the approximated kernel space by the Nyström kernel expansion (see e.g. Williams and Seeger (2000)).

For the ball data set the data contain substantial information in the negative fraction of the eigenspectrum, accordingly one may expect that these eigenvalues should not be removed. This is also reflected in the results. In SIM4 and SIM 5 only the two dominating eigenvectors are kept such that the negative eigenvalues are removed, degenerating the prediction accuracy. The SIM3 encoding is a bit better, but the landmark optimization via k-means is not very effective for this dataset. Also the entropy approach in SIM6 was not very efficient. The SIM2 encoding has a substantial drop in the accuracy with respect to the unapproximated kernel but the intrinsic dimension of the dataset is very high and the $m = 8$ landmarks are enough to preserve the dominating positive and negative eigenvalues. The unapproximated kernel leads to perfect separation, clearly showing that the negative eigenspectrum contains discriminative information. The respective eigenvalue plots are provided in Figure 3.

The results show that the proposed MEB approach is capable in preserving the geometric information also for the negative (squared) eigendimensions while being quite simple. We believe that controlling the approximation accuracy of the kernel by ϵ in the MEB is much easier than selecting the number of clusters (per class) in k-means clustering. In fact it will almost always be sufficient to keep $\epsilon \approx 0.01$ to get reliable landmark sets whereas the number of clusters is very dataset dependent and not easy to choose. However, in contrast to the results

[9] We use the implementation as provided by the authors in the LSSVM toolbox `http://www.esat.kuleuven.be/sista/lssvmlab/`

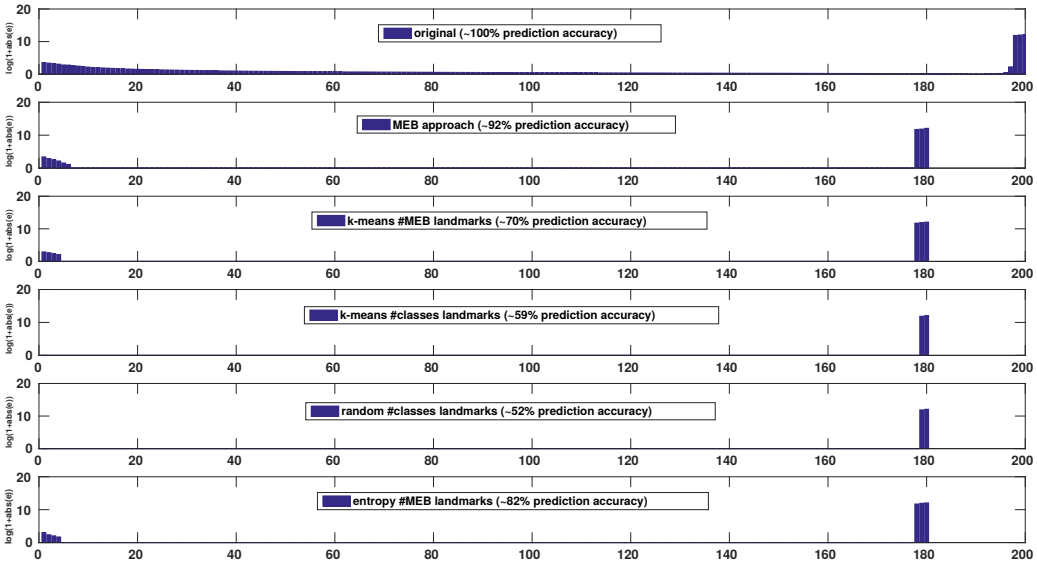
original (~100% prediction accuracy)
MEB approach (~92% prediction accuracy)
k-means #MEB landmarks (~70% prediction accuracy)
k-means #classes landmarks (~59% prediction accuracy)
random #classes landmarks (~52% prediction accuracy)
entropy #MEB landmarks (~82% prediction accuracy)

Figure 3: Eigenvalue analysis of the ball dataset using the different approaches. The first plot shows the eigenvalues of the original kernel (SIM1), the other plots show typical results from the 10-fold crossvalidation for the various landmark selection approaches (SIM2-SIM6). It can be clearly seen that the landmarks identified by the MEB approach sufficiently capture the negative eigenvalues. The random sampling approach works only if a larger number of landmarks is chosen and is still less efficient because it is not ensured that the landmarks cover the whole data space. Especially if the data are non i.i.d. random sampling is typically insufficient.

Table 1: Test set results of a 10-fold iKFD run on the simulated / controlled datasets in different kernel approximations and the obtained SMSS (median) value. A $\star$ indicates a non-metric similarity matrix. The number of identified landmarks is shown in brackets for SIM2.

Method	Ball$\star$	Protein$\star$	Checker	Gaussian
$SIM1\|s(\hat{K},K)$	$100 \pm 0\|1$	$98.12 \pm 3.22\|1.0$	$98.89 \pm 0.35\|1.0$	$90.00 \pm 5.77\|1.0$
$SIM2\|s(\hat{K},K)$	$92.00 \pm 4.83\|1.00(8)$	$96.71 \pm 3.20\|0.98(25)$	$90.22 \pm 8.52\|0.78(9)$	$90.00 \pm 7.45\|0.93(8)$
$SIM3\|s(\hat{K},K)$	$70.00 \pm 12.69\|0.99$	$96.71 \pm 4.45\|0.85$	$91.78 \pm 9.24\|1.0$	$87.00 \pm 10.33\|0.98$
$SIM4\|s(\hat{K},K)$	$59.50 \pm 5.50\|0.40$	$86.85 \pm 6.29\|0.71$	$65.33 \pm 5.13\|0.04$	$65.00 \pm 8.17\|0.13$
$SIM5\|s(\hat{K},K)$	$52.50 \pm 12.08\|0.37$	$78.87 \pm 14.61\|1.01$	$46.11 \pm 4.20\|0.18$	$77.50 \pm 10.61\|0.40$
$SIM6\|s(\hat{K},K)$	$74.50 \pm 12.79\|1.00$	$95.31 \pm 5.78\|0.81$	$62.33 \pm 11.67\|0.10$	$87.00 \pm 7.52\|0.36$

Table 2: Runtimes of a 10-fold iKFD run on the simulated / controlled datasets in different kernel approximations. A $\star$ indicates a non-metric similarity matrix.

Method	Ball$\star$	Protein$\star$	Checker	Gaussian
$SIM1\|s(\hat{K},K)$	0.5	0.82	13.45	0.74
$SIM2\|s(\hat{K},K)$	1.0	1.56	3.76	0.98
$SIM3\|s(\hat{K},K)$	1.57	2.57	14.77	1.51
$SIM4\|s(\hat{K},K)$	0.84	1.14	13.23	0.90
$SIM5\|s(\hat{K},K)$	0.61	0.98	3.23	0.65
$SIM6\|s(\hat{K},K)$	3.2	8.47	8.12	3.94

shown in Table 1 the approach by Zhang and Kwok (2010) is typically effective for a large variety of datasets also with indefinite kernels, given the number of landmarks is reasonable large and discriminating information is sufficiently provided in the dominating eigenvectors of the cluster solutions. For the protein data we observe similar results and the proposed approach, the k-means strategy and the entropy approach are effective. SIM4 and SIM5 is again substantially worse because four landmarks are in general not sufficient to represent these data from a discriminative point of view.

For the checker board and Gaussian data SIM2 and SIM3 are again close and SIM4 and SIM5 are substantially worse using only two landmark points. The entropy approach was efficient only for the Gaussian data, but failed for Checker which may be attributed to the strong multi-modality of the data.

The runtimes shown in Table 2 show already for the small data examples that the MEB approach is much faster then k-means or the entropy approach if the number of points gets larger which was already expected from the theoretical runtime complexity of these algorithms.

In another small experiment we analyzed the effect of the k-means based landmark selection Zhang and Kwok (2010) in more detail. We consider three Gaus-

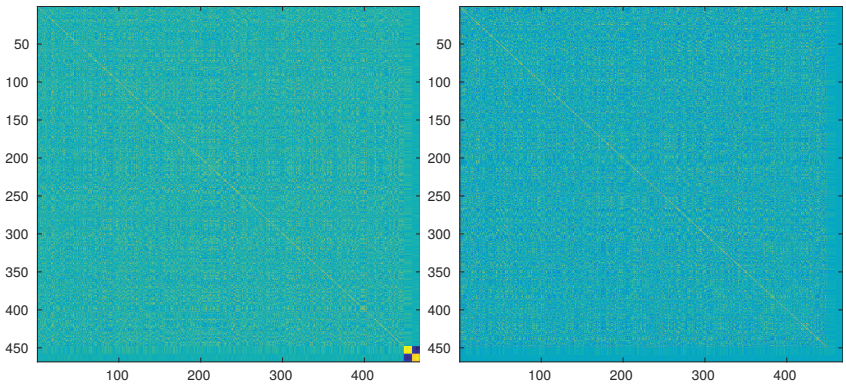

Figure 4: Reconstructed kernel matrix (from the crossvalidation run) of the 10 dimensional Gaussian example. Left using the MEB approach, right using the k-means landmark selection. Note the small region on the bottom in the left plot indicating the smaller gaussians which are almost missing in the right plot.

sians where one Gaussian has 500 points spread in two dimensions and two other Gaussians each with 20 points spread in another dimensions. All Gaussians are perfectly separated to each other located in a three dimensional space. To make the task more challenging we further add 7 dimensions with small noise contributions to the large Gaussian. The final data are given in a 10 dimensional space, whereby the small Gaussians are intrinsically low dimensional and the large Gaussian is 10 dimensional. with major contributions only in two dimensions. The points from the large Gaussian are labeled 0 and the other 1. Using the MEB approach we obtain 10 landmarks and the approximated kernel is sufficient to give a perfect prediction of 100% in a 10-fold crossvalidation with iKFD. Using the k-means or entropy based approach (with the same number of landmarks) the prediction accuracy drops down to $\approx 84\%$ and for random sampling we get a prediction accuracy in the same range of 83% - again with 10 landmarks . This can be explained by the behavior of k-means to assign the prototypes or landmarks to dense regions. It is hence more likely that after the k-means clustering (almost) all prototypes are used to represent the large Gaussian and no prototypes are left for the other classes. Due to the fact that the other classes are located in different dimensions with respect to the large Gaussian these dimensions are not any longer well represented and hence the respective classes are often missing in the approximated kernel (see Figure 4). This density related behavior is also known as magnification Villmann, T., Claussen, J. C., (2006) in the context of different vector quantization approaches. Hence using the unsupervised k-means landmark selection it

Algorithm 1 Proposed handling of indefinite kernels by the MEB approach

1. let $k(x, y)$ be a symmetric (indefinite) similarity function (e.g. a sequence alignment)
2. for all labels c let $D_c = \{(x_i, y_i) : y_i = c\}$
3. calculate the (indefinite) kernel matrix K_c using D_c and $k(x, y)$
4. if the kernel matrix is indefinite, apply a square operation on the small matrix K_c by using $K_c \cdot K_c^\top$
5. calculate $\forall x \Delta(x) = k(x, x)$ (respectively for all K_c)
6. apply the MEB algorithm for each of the kernel matrices K_c with $\epsilon = 0.01$ and the respective subset of Δ
7. combine all landmark indices obtained from the former step and calculate the Nyström approximation using Eq. (9)
8. apply Ny-PCVM or Ny-iKFD using the approximated kernel matrix

can easily happen, that the majority of the data space is well presented but small classes are ignored - which is obviously a problem for a supervised data analysis.

From these initial experiments we see that the proposed landmark selection scheme is sufficient to approximate the original kernel function for a *supervised* analysis as indicated by the prediction accuracy of the iKFD model and the SMSS value. We also see that the Nyström approximation can introduce substantial error if the data are not low rank (for checker) due to a more complicated kernel mapping aka similarity function. We would like to highlight again that without an advocated guess of the number of landmarks neither the k-means strategy nor the entropy approach are very efficient.

In the experiment in section 7 we will restrict our analysis to the proposed landmark selection using the MEB approach, the k-means strategy and the entropy based technique.

5. Large scale indefinite learning with PCVM and iKFD

We now integrate the aforementioned Nyström approximation approaches and the supervised landmark selection into PCVM and iKFD. The modifications ensure that all matrices are processed with linear memory complexity and that the underlying algorithms have a linear runtime complexity. For both algorithms the initial input is the Nyström approximated kernel matrix with landmarks selected by using one of the formerly provided landmark selection schemes.

5.1. PCVM for large scale proximity data

The PCVM parameters are optimized using the EM algorithm to prune the weight vector $\mathbf{w}$ during learning and hence the considered basis functions representing the model. We will now show multiple modifications of PCVM to integrate the Nyström approximation and to ensure that the memory and runtime complexity remains linear at all time. We refer to our method as Ny-PCVM. Initially the Ny-PCVM algorithm makes use of the matrices $K_1 = K_{N,m}$ and $K_2 = K_{m,m}^{-1} \cdot K_1^\top$ obtained from the original kernel matrix using the Nyström landmark technique described above. Given a matrix X, we denote by $\hat{X}$ the matrix formed from X containing elements at indices that have not yet been pruned out of the weight vector $\mathbf{w}$. As an example, the matrices $\hat{K_1} = K_1^{\mathbf{w}\neq 0,\cdot}$, $\hat{K_2} = K_2^{\cdot,\mathbf{w}\neq 0}$ hold only those columns/rows of K_1 or K_2 not yet pruned out from the weight vector. We will use the same notation also for other variables. We denote the set of indices of m randomly selected landmarks by $[m]$. Finally, in contrast to the original PCVM formulation Chen et al. (2009a), in our notation we explicitly use the data labels - for example, instead of vector $\Phi_\theta(\mathbf{x})$ we write $\Xi_\theta(\mathbf{x}) \circ \mathbf{y}$, where $\Xi_\theta(\mathbf{x})$ is the kernel vector of $\mathbf{x}$ without any label information, $\mathbf{y}$ is the label vector and $\circ$ is the element-wise multiplication.

We now adapt multiple equations of the original PCVM to integrate the Nyström approximated matrix. Beginning with the elements of vector (for a *single* training vector i) $\mathbf{z}_\theta$:

$$z_{i,\theta} = \Xi_\theta(\mathbf{x}_i)(\mathbf{y} \circ \mathbf{w}) + b, \tag{10}$$

we rewrite Eq.(10) in matrix notation for all training points:

$$\hat{\mathbf{z}} = (((\hat{\mathbf{y}} \circ \hat{\mathbf{w}})^\top \hat{K_1}) \cdot K_2)^\top + b \tag{11}$$

and further obtain column vectors $\bar{\mathbf{H}}_\theta$ and the reduced form $\bar{\hat{\mathbf{H}}}_\theta$, by using only the non-vanishing basis functions and the Nyström approximated matrices in Eq. (4). In the maximization step of the original PCVM the $\mathbf{w}$ are updated as (see Eq. (5)):

$$\mathbf{w}^{\text{new}} = M\underbrace{(M\Phi_\theta(\mathbf{x})^\top \Phi_\theta(\mathbf{x})M + I_N)}_{\Upsilon}{}^{-1} M(\Phi_\theta(\mathbf{x})^\top \bar{\mathbf{H}}_\theta - b\Phi_\theta(\mathbf{x})^\top \mathbf{I}) \tag{12}$$

To account for the now excluded labels we reformulate Equation (5) as:

$$\mathbf{w}^{\text{new}} = M\underbrace{(M(\Xi_\theta(\mathbf{x})^\top \Xi_\theta(\mathbf{x})\hat{\mathbf{y}}^\top \hat{\mathbf{y}})M + I_N)}_{\Upsilon}{}^{-1} M(\hat{\mathbf{y}}^\top(\Xi_\theta(\mathbf{x})^\top \bar{\mathbf{H}}_\theta) - b\hat{\mathbf{y}}^\top(\Xi_\theta(\mathbf{x})^\top \mathbf{I}))$$

The update equations of the weight vector include the calculation of a matrix inverse of Υ which was originally calculated using the Cholesky decomposition. To keep our objective of small matrices we will instead calculate the pseudo-inverse of this matrix using a Nyström approximation of Υ. It should be noted at this point that the matrix Υ is psd by construction. We approximate Υ by selecting another set of m^* landmarks from the indices of the not yet pruned weights and calculate the matrix $\tilde{\Upsilon} = C_{Nm^*} W^{-1}_{m^*,m^*} C^\top_{Nm^*}$ in analogy to Eq (9) with submatrices: [10]

$$\begin{aligned} C_{Nm^*} &= E_{N[m]} + ((\hat{K_1} \cdot (K_2 \cdot (K_1 \cdot \hat{K}_{2\cdot,[m^*]}))(\hat{\mathbf{y}}^\top \hat{\mathbf{y}}_{[m^*]})) \\ &\quad \circ \sqrt{2}\hat{\mathbf{w}}) \circ \sqrt{2}\hat{\mathbf{w}}^\top_{[m^*]} \\ W_{m^*,m^*} &= C^{-1}_{m^*,\cdot} \end{aligned}$$

Where $\circ$ indicates (in analogy to its previous meaning) that each row of the left matrix is elementwise multiplied by the right vector and $E_{N[m]}$ is the matrix consisting of the m landmark columns of the $N \times N$ identity matrix. The terms $\sqrt{2}\hat{\mathbf{w}}$ and $\sqrt{2}\hat{\mathbf{w}}^\top_{[m^*]}$ are the entries of the diagonal matrix M as defined in Eq. (7) but now given in vector form.

These two matrices serve as the input of a Nyström approximation based pseudo-inverse (as discussed in sub section 3.1) and we obtain matrices $V \in \mathbb{R}^{N\times r}, U \in \mathbb{R}^{r\times N}$ and $S \in \mathbb{R}^{r\times r}$, where $r \leq m^*$ is the rank of the pseudo inverse. Further we define two vectors $\mathbf{v}_1 = \bar{\hat{\mathbf{H}}}_\theta^\top \cdot K_1$ and $\mathbf{v}_2 = \mathbf{I}^\top \cdot K_1$. We obtain the approximated weight update $\mathbf{w}^{\text{new}} = V \cdot (S \cdot U^\top \cdot (\sqrt{2}\hat{\mathbf{w}}(\hat{\mathbf{y}}(\mathbf{v}_1 \cdot \hat{K_2})^\top - b \cdot \hat{\mathbf{y}}(\mathbf{v}_2 \cdot \hat{K_2})^\top)))\sqrt{2}\hat{\mathbf{w}}$. The update of the bias is originally done as

$$\mathbf{b} = t(1 + tNt)^{-1} t(\mathbf{I}^\top \bar{\hat{\mathbf{H}}}_\theta - \mathbf{I}^\top \Phi_\theta(\hat{\mathbf{y}}\hat{\mathbf{w}})) \tag{13}$$

which is replaced to: $\mathbf{b} = t(1 + tNt)^{-1} t(\mathbf{I}^\top \bar{\hat{\mathbf{H}}}_\theta - \mathbf{I}^\top ((((\hat{\mathbf{y}}\hat{\mathbf{w}})^\top \hat{K_1}) \cdot K_2)^\top))$ Subsequently the entries in $\hat{\mathbf{w}}$ which are close to zero are pruned out and the matrices $\hat{K_1}$ and $\hat{K_2}$ are modified accordingly.

5.2. Nyström based Indefinite Kernel Fisher Discriminant

Given a Nyström approximated kernel matrix a few adaptations have to be made to obtain a valid iKFD formulation solely based on the Nyström approximated kernel, without any full matrix operations.

[10] The number of landmarks m^* is fixed to be 1% of $|w|$ but not more then 500 landmarks. If the length of $\mathbf{w}$ drops below 100 points we use the original PCVM formulations.

First we need to calculate the classwise means μ_+ and μ_- based on the row/column sums of the approximated input kernel matrix. This can be done by rather simple matrix operations on the two low rank matrices of the Nyström approximation of K. For better notation let us define the matrices K_{Nm} as Ψ and K_{mm} as Γ then for each row k of the matrix K we get the row/column sum as:

$$\sum_i [\tilde{K}]_{k,i} = \sum_{l=1}^{m} \left(\sum_{j=1}^{N} \Psi_{j,\cdot} \Gamma^{-1} \right) \Psi_{l,k}^{\top} \tag{14}$$

This can obviously also be done in a single matrix operation for all rows in a batch, with linear complexity only. Based on these mean estimates we can calculate Eq. (2). In a next step we need to calculate a squared approximated kernel matrix for the positive and the negative class with removed means μ_+ or μ_- respectively. For the positive class with n_+ entries, we can define a new Nyström approximated (squared) matrix with subtracted mean as :

$$\hat{K}^{+}_{N,m} = K_{N,m} \cdot K^{-1}_{m,m} \cdot (K^{\top}_{I_+,m} \cdot K_{I_+,m}) \cdot K^{-1}_{m,m} \cdot K^{\top}_{m,m} - \mu_+ \cdot \mu_+^{\top} \cdot n_+ \tag{15}$$

An equivalent term can be derived for the negative class providing $\hat{K}^{-}_{N,m}$. It should be noted that no obtained matrix in Eq (15) has more than $N \times m$ entries. Finally $\hat{K}^{+}_{N,m}$ and $\hat{K}^{-}_{N,m}$ are combined to approximate the within class matrix as shown in Eq. (3). From the derivation in Haasdonk and Pekalska (2008) we know, that only the eigenvector of the Nyström approximated kernel matrix based on $\hat{K}_{N,m} = \hat{K}^{+}_{N,m} + \hat{K}^{-}_{N,m}$ are needed. Using a Nyström based eigen-decomposition (explained before) on $\hat{K}_{N,m}$ we obtain:

$$\alpha = C \cdot A^{-1} \cdot (C' \cdot (\mu_+ - \mu_-))$$

where C contains the eigenvectors and A the eigenvalues of $\hat{K}_{N,m}$. Instead of A^{-1} one can use the pseudo-inverse. The bias term b is obtained as $b = -\alpha^{\top}(\mu_+ + \mu_-)/2$.

6. Complexity analysis

The original iKFD update rules have costs of $\mathcal{O}(N^3)$ and memory storage $\mathcal{O}(N^2)$, where N is the number of points. The Ny-iKFD may involve the extra Nyström approximation of the kernel matrix to obtain $K_{N,m}$ and $K^{-1}_{m,m}$, if not already given. If we have m landmarks, $m \ll N$, this gives costs of $\mathcal{O}(mN)$ for the

first matrix and $\mathcal{O}(m^3)$ for the second, due to the matrix inversion. Further both matrices are multiplied within the optimization so we get $\mathcal{O}(m^2N)$. Similarly, the matrix inversion of the original iKFD with $\mathcal{O}(N^3)$ is reduced to $\mathcal{O}(m^2N)+\mathcal{O}(m^3)$ due to the Nyström approximation of the pseudo-inverse. If we assume $m \ll N$ the overall runtime and memory complexity of Ny-iKFD is linear in N. For the Ny-PCVM we obtain a similar analysis as shown in Schleif et al. (2015a) but with extra costs to calculate the Nyström approximated SVD. Additionally, Ny-PCVM uses an iterative optimization scheme to optimize and sparsify w with constant costs C_I, as the number of iterations. Accordingly Ny-iKFD and Ny-PCVM have both linear memory and runtime complexity $\mathcal{O}(N)$, but Ny-PCVM maybe slower than Ny-iKFD due to extra overhead costs.

7. Experiments

We compare iKFD, Ny-iKFD, Ny-PCVM and PCVM on various larger indefinite proximity data. In contrast to many standard kernel approaches, for iKFD and PCVM, the indefinite kernel matrices need not to be corrected by costly eigenvalue correction Chen et al. (2009c); Schleif and Gisbrecht (2013) [11]

Further the iKFD and PCVM provides direct access to probabilistic classification decisions. First we show a small simulated experiment for two Gaussians which exist in an intrinsically two dimensional *pseudo*-Euclidean space $\mathbb{R}^{(1,1)}$. The plot in Figure 5 shows a typical result for the obtained decision planes using the iKFD or Ny-iKFD. The Gaussians are slightly overlapping and both approaches achieve a good separation with 93.50% and 88.50% prediction accuracy, respectively.

Subsequently we consider a few public available datasets for some real life experiments. The data are *Zongker* (2000pts, 10 classes) and *Proteom* (2604pts, 53 classes (restricted to classes with at least 10 entries)) from Duin (2012); *Chromo* (4200pt, 21 classes) from Neuhaus and Bunke (2006) and the SwissProt database *Swiss* (10988 pts, 30 classes) from Boeckmann et al. (2003), (version 10/2010, reduced to prosite labeled classes with at least 100 entries). Further we used the *Sonatas* data (1068pts, 5 classes) taken from Mokbel et al. (2009). All data are processed as indefinite kernels and the landmarks are selected using the respective landmark selection schemes. The mean number of Nyström landmarks as obtained by the MEB approach is given in brackets after the dataset label. For all

[11] In Schleif and Tino (2015) various correction methods have been studied on the same data indicating that eigenvalue corrections may be helpful.

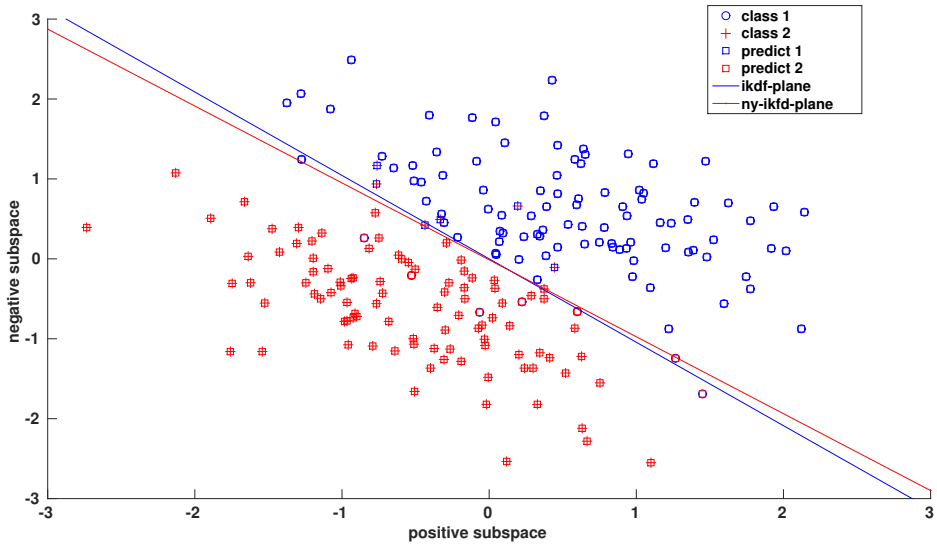

Figure 5: Visualization of the indefinite Fisher kernel for two Gaussians in a two dimensional pseudo-Euclidean space $\mathbb{R}^{(1,1)}$. The predicted labels are with respect to the iKFD classification.

experiments we report mean and standard errors as obtained by a 10 fold crossvalidation. For PCVM we fixed the upper number of optimization cycles to 500. The probabilistic outputs can be directly used to allow for a reject region but can also be used to provide alternative classification decisions e.g. in a ranking framework

In Table 4, 3 and Table 6 we show the results for different non-metric proximity datasets using Ny-PCVM, PCVM and iKFD or Ny-iKFD. The overall best results for a dataset are underlined and the best approximations are highlighted in bold. Considering Table 4 and Table 3 we see that iKFD and PCVM are similarly effective with slightly better results for iKFD. The Nyström approximation of the kernel matrix *only*, often leads to a in general small decrease of the accuracy, but the additional approximation step, in the algorithm itself, does not substantially decrease the prediction accuracy further[12].

Considering the overall results in Table 4 and Table 3 the approximations used in the algorithms Ny-iKFD and Ny-PCVM appear to be effective. The runtime analysis in Table 6 clearly shows that the classical iKFD is very complex. As ex-

[12] Also the runtime and model complexity are similar and therefore not reported in the following.

dataset	iKFD	Ny-iKFD (MEB)	PCVM	Ny-PCVM (MEB)	$\overline{SMSS}$
gesture (64)	$\underline{97.93} \pm 0.73$	$\mathbf{96.60 \pm 1.84}$	73.20 ± 18.12	$85.53 \pm 1.22*$	0.9481
sonatas (25)	90.17 ± 2.14	$83.52 \pm 2.08*$	$\underline{91.20} \pm 2.69$	$\mathbf{87.08 \pm 3.19}*$	0.7460
zongker (41)	$\underline{96.60} \pm 1.97$	$\mathbf{90.70 \pm 2.30}*$	93.60 ± 2.00	$84.35 \pm 2.53*$	0.6785
proteom (123)	99.58 ± 0.38	$\mathbf{\underline{99.68} \pm 0.31}$	99.58 ± 0.28	99.45 ± 0.53	0.9184
chromo (65)	$\underline{97.24} \pm 0.94$	$\mathbf{94.79 \pm 1.45}$	93.29 ± 1.51	92.21 ± 1.31	0.9300
swiss (116)	–	$\underline{83.05} \pm 1.60$	–	70.38 ± 19.19	0.7870

Table 3: Crossvalidation results using the MEB approach. iKFD and PCVM use the original indefinite kernel without approximations. Ny-iKFD and Ny-PCVM use the Nyström approximation within the implementation as discussed before and the same Nyström approximated kernel. (*) indicate significant differences with respect to the same unapproximated method. The mean $\overline{SMSS}$ values are calculated on the MEB based Nyström approximation.

pected, the integration of the Nyström approximation leads to substantial speed-ups. Larger datasets like the Swiss data with ≈ 10.000 entries could not be analyzed by iKFD or PCVM before. We also see that the landmark selection scheme using MEB is slightly more effective than by using k-means but without the need to tune the number of clusters (landmarks). The entropy approach is similar efficient than the k-means strategy but more costly due to the iterative optimization of the landmark set and the respective eigen-decompositions (see Brabanter et al. (2010)).

The PCVM is focusing on a sparse parameter vector w in contrast to the iKFD. For the iKFD most training points are also used in the model ($\geq 94\%$) whereas for Ny-PCVM often less than 5% are kept in general as shown in Table 7. In practice it is often costly to calculate the non-metric proximity measures like sequence alignments and also a large number of kernel expansions should be avoided. Accordingly sparse models are very desirable. Considering the runtime again Ny-PCVM and Ny-iKFD are in general faster than the original algorithms, typically by at least a magnitude. the PCVM and Ny-PCVM are also very fast in the test case or out-of sample extension due to the inherent model sparsity.

8. Conclusions

We presented an alternative formulation of the iKFD and PCVM employing the Nyström approximation. We also provided an alternative way to identify the landmark points of the Nyström approximation in cases where the objective is

dataset	iKFD	Ny-iKFD (KM)	PCVM	Ny-PCVM (KM)	$\overline{SMSS}$
gesture (64)	$\underline{97.93} \pm 0.73$	95.73 ± 0.86	73.20 ± 18.12	$92.60 \pm 1.04*$	0.9315
sonatas (25)	90.17 ± 2.14	$77.63 \pm 3.19*$	$\underline{91.20} \pm 2.69$	$77.81 \pm 3.28*$	0.4925
zongker (41)	$\underline{96.60} \pm 1.97$	$88.40 \pm 1.33*$	93.60 ± 2.00	$88.30 \pm 2.89*$	0.8340
proteom (123)	99.58 ± 0.38	94.78 ± 1.89	99.58 ± 0.28	94.18 ± 1.23	0.5711
chromo (65)	$\underline{97.24} \pm 0.94$	94.17 ± 0.86	93.29 ± 1.51	92.10 ± 0.89	0.9406
swiss (116)	–	73.74 ± 0.71	–	75.36 ± 7.55	0.7864

Table 4: Crossvalidation results using the k-means (KM) approach. iKFD and PCVM use the original indefinite kernel without approximations. Ny-iKFD and Ny-PCVM are Nyström approximated and use the same Nyström approximated kernel obtained by the KM strategy. (*) indicate significant differences with respect to the same unapproximated method. The number of landmarks is chosen w.r.t. the MEB solution. The mean $\overline{SMSS}$ values are calculated on the KM based Nyström approximation.

a supervised problem. Our results indicate that in general the MEB approach is similar efficient compared to the k-means clustering or the entropy strategy but with less effort and almost parameter free. We found that Ny-iKFD is competitive in the prediction accuracy with the original iKFD and alternative approaches, while taking substantially less memory and runtime but being less sparse then Ny-PCVM. The Ny-iKFD and Ny-PCVM provides now an effective way to obtain a *probabilistic* classification model for medium to large psd *and* non-psd datasets, in *batch mode* with *linear* runtime and memory complexity. If sparsity is not an issue one may prefer Ny-iKFD which is slightly better in the prediction accuracy then Ny-PCVM. Using the presented approach we believe that iKFD is now applicable for realistic problems and may get a larger impact then before. In future work it could be interesting to incorporate sparsity concepts into iKFD and Ny-iKFD similar as shown for classical KFD in Diethe et al. (2009).

Implementation: The Nyström approximation for iKFD is provided at `http://www.techfak.uni-bielefeld.de/~fschleif/source/ny_ikfd.tgz` and the PCVM/Ny-PCVM code can be found at `https://mloss.org/software/view/610/`.

Acknowledgment: A Marie Curie Intra-European Fellowship (IEF): FP7-PEOPLE-

dataset	iKFD	Ny-iKFD (ENT)	PCVM	Ny-PCVM (ENT)	$\overline{SMSS}$
gesture (64)	97.93 ± 0.73	93.47 ± 1.93*	73.20 ± 18.12	91.07 ± 2.97*	0.8864
sonatas (25)	90.17 ± 2.14	80.24 ± 2.46*	91.20 ± 2.69	82.77 ± 2.86*	0.5210
zongker (41)	96.60 ± 1.97	90.90 ± 1.15*	93.60 ± 2.00	90.50 ± 2.12	0.5500
proteom (123)	99.58 ± 0.38	94.54 ± 1.87	99.58 ± 0.28	80.93 ± 22.96*	0.5057
chromo (65)	97.24 ± 0.94	94.50 ± 1.30	93.29 ± 1.51	90.95 ± 2.55	0.9560
swiss (116)	-	-	-	-	

Table 5: Crossvalidation results using the entropy (ENT) approach. Ny-iKFD and Ny-PCVM are Nyström approximated and use the same Nyström approximated kernel obtained by the entropy strategy. (*) indicate significant differences with respect to the same unapproximated method. The number of landmarks is chosen w.r.t. the MEB solution. The mean $\overline{SMSS}$ values are calculated on the ENT based Nyström approximation. For the swiss data the entropy approach took to much time.

	iKFD	Ny-iKFD	PCVM	Ny-PCVM
gesture	50.72 ± 1.54	9.18 ± 0.19	116.33 ± 7.49	31.98 ± 0.42
sonatas	5.04 ± 0.22	1.85 ± 0.06	60.07 ± 2.54	7.01 ± 0.24
zongker	51.61 ± 1.43	5.53 ± 0.16	184.07 ± 14.97	16.91 ± 0.24
proteom	559.25 ± 15.29	42.08 ± 1.92	352.08 ± 18.05	111.22 ± 1.88
chromo	763.24 ± 31.54	27.91 ± 1.77	694.43 ± 15.61	54.36 ± 0.77
swiss	–	178.79 ± 10.63	–	123.29 ± 2.72

Table 6: Typical runtimes - indefinite kernels

2012-IEF (FP7-327791-ProMoS) and support from the Cluster of Excellence 277 Cognitive Interaction Technology funded by the German Excellence Initiative is gratefully acknowledged. PT was supported by the EPSRC grant EP/L000296/1, "Personalized Health Care through Learning in the Model Space". We would like to thank R. Duin, Delft University for various support with distools and prtools and Huanhuan Chen,University of Science and Technology of China, for providing support with the Probabilistic Classification Vector Machine.

References

Alabdulmohsin, I. M., Gao, X., Zhang, X., 2014. Support vector machines with indefinite kernels. In: Phung, D., Li, H. (Eds.), Proceedings of the Sixth Asian

	iKFD	Ny-iKFD (MEB)	PCVM	Ny-PCVM (MEB)
gesture	100.00 ± 0	100.00 ± 0	10.60 ± 0.84	5.25 ± 0.31
sonatas	100.00 ± 0	100.00 ± 0	11.24 ± 0.56	3.42 ± 0.57
zongker	100.00 ± 0	100.00 ± 0	14.42 ± 3.65	8.63 ± 0.31
proteom	100.00 ± 0	100.00 ± 0	5.23 ± 0.36	5.85 ± 0.14
chromo	100.00 ± 0	100.00 ± 0	7.49 ± 0.51	2.49 ± 0.34
swiss	$-$	96.95 ± 0.27	$-$	1.18 ± 0.25

Table 7: Model complexity - indefinite kernels (threshold $1e^{-4}$)

Conference on Machine Learning, ACML 2014, Nha Trang City, Vietnam, November 26-28, 2014. Vol. 39 of JMLR Proceedings. JMLR.org, p. in press.

Badoiu, M., Clarkson, K. L., 2008. Optimal core-sets for balls. Comput. Geom. 40 (1), 14–22.

Balcan, M.-F., Blum, A., Srebro, N., 2008. A theory of learning with similarity functions. Machine Learning 72 (1-2), 89–112.

Belkin, M., Niyogi, P., 2003. Laplacian eigenmaps for dimensionality reduction and data representation. Neural Computation 15 (6), 1373–1396.
URL http://dx.doi.org/10.1162/089976603321780317

Boeckmann, B., Bairoch, A., Apweiler, R., Blatter, M.-C., Estreicher, A., Gasteiger, E., Martin, M., Michoud, K., O'Donovan, C., Phan, I., Pilbout, S., Schneider, M., 2003. The SWISS-PROT protein knowledgebase and its supplement TrEMBL in 2003,. Nucleic Acids Research 31, 365–370.

Brabanter, K. D., Brabanter, J. D., Suykens, J., Moor, B. D., 2010. Optimized fixed-size kernel models for large data sets. Computational Statistics & Data Analysis 54 (6), 1484 1504.

Chen, H., Tino, P., Yao, X., 2009a. Probabilistic classification vector machines. IEEE Transactions on Neural Networks 20 (6), 901–914.

Chen, H., Tino, P., Yao, X., 2014. Efficient probabilistic classification vector machine with incremental basis function selection. IEEE TNN-LS 25 (2), 356–369.

Chen, Y., Garcia, E. K., Gupta, M. R., Rahimi, A., Cazzanti, L., 2009c. Similarity-based classification: Concepts and algorithms. JMLR 10, 747–776.

Cichocki, A., Amari, S.-I., 2010. Families of alpha- beta- and gamma- divergences: Flexible and robust measures of similarities. Entropy 12 (6), 1532–1568.

Diethe, T., Hussain, Z., Hardoon, D. R., Shawe-Taylor, J., 2009. Matching pursuit kernel fisher discriminant analysis. In: Dyk, D. A. V., Welling, M. (Eds.), Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics, AISTATS 2009, Clearwater Beach, Florida, USA, April 16-18, 2009. Vol. 5 of JMLR Proceedings. JMLR.org, pp. 121–128.

Dubuisson, M., Jain, A., Oct 1994. A modified hausdorff distance for object matching. In: Pattern Recognition, 1994. Vol. 1 - Conference A: Computer Vision amp; Image Processing., Proceedings of the 12th IAPR International Conference on. Vol. 1. pp. 566–568 vol.1.

Duin, R. P., march 2012. prtools.
URL `http://www.prtools.org`

Duin, R. P. W., Pekalska, E., 2010. Non-euclidean dissimilarities: Causes and informativeness. In: Structural, Syntactic, and Statistical Pattern Recognition, Joint IAPR International Workshop, SSPR&SPR 2010, Cesme, Izmir, Turkey, August 18-20, 2010. Proceedings. pp. 324–333.

Gisbrecht, A., Schleif, F.-M., 2015. Metric and non-metric proximity transformations at linear costs. Neurocomputing to appear.

Gittens, A., Mahoney, M. W., 2013. Revisiting the nyström method for improved large-scale machine learning. CoRR abs/1303.1849.

Goldfarb, L., 1984. A unified approach to pattern recognition. Pattern Recognition 17 (5), 575 – 582.

Haasdonk, B., 2005. Feature space interpretation of svms with indefinite kernels. IEEE Transactions on Pattern Analysis and Machine Intelligence 27 (4), 482–492.

Haasdonk, B., Keysers, D., 2002. Tangent distance kernels for support vector machines. In: ICPR (2). pp. 864–868.

Haasdonk, B., Pekalska, E., 2008. Indefinite kernel fisher discriminant. In: 19th International Conference on Pattern Recognition (ICPR 2008), December 8-11, 2008, Tampa, Florida, USA. IEEE Computer Society, pp. 1–4.

Kwok, J. T., Tsang, I. W., 2003. Learning with idealized kernels. In: Fawcett, T., Mishra, N. (Eds.), Machine Learning, Proceedings of the Twentieth International Conference (ICML 2003), August 21-24, 2003, Washington, DC, USA. AAAI Press, pp. 400–407.
URL http://www.aaai.org/Library/ICML/2003/icml03-054.php

Ling, H., Jacobs, D. W., 2007. Shape classification using the inner-distance. IEEE Trans. Pattern Anal. Mach. Intell. 29 (2), 286–299.

Loosli, G., Canu, S., Ong, C., 2015. Learning svm in krein spaces. Pattern Analysis and Machine Intelligence, IEEE Transactions on PP (99), 11.

Mokbel, B., Hasenfuss, A., Hammer, B., 2009. Graph-based representation of symbolic musical data. In: Torsello, A., Escolano, F., Brun, L. (Eds.), Graph-Based Representations in Pattern Recognition, 7th IAPR-TC-15 International Workshop, GbRPR 2009, Venice, Italy, May 26-28, 2009. Proceedings. Vol. 5534 of Lecture Notes in Computer Science. Springer, pp. 42–51.

Nebel, D., Hammer, B., Frohberg, K., Villmann, T., 2015. Median variants of learning vector quantization for learning of dissimilarity data. NeurocomputingCited By 0; Article in Press.
URL http://www.scopus.com/inward/record.url?eid=2-s2.0-84930918787&partnerID=40&md5=43de27cecb1b43b7769b0baf6e0744a3

Neuhaus, M., Bunke, H., 2006. Edit distance based kernel functions for structural pattern classification. Pattern Recognition 39 (10), 1852–1863.

Pekalska, E., Duin, R., 2005. The dissimilarity representation for pattern recognition. World Scientific.

Pekalska, E., Haasdonk, B., 2009. Kernel discriminant analysis for positive definite and indefinite kernels. IEEE Transactions on Pattern Analysis and Machine Intelligence 31 (6), 1017–1031.

Schleif, F.-M., A.Gisbrecht, Tino, P., 2015. Probabilistic classification vector machine at large scale. In: Proceedings of ESANN 2015. pp. 555–560.

Schleif, F.-M., Gisbrecht, A., 2013. Data analysis of (non-)metric proximities at linear costs. In: Proceedings of SIMBAD 2013. pp. 59–74.

Schleif, F.-M., Gisbrecht, A., Tino, P., 2015b. Large scale indefinite kernel fisher discriminant. In: Proceedings of Simbad 2015. pp. 160170.

Schleif, F.-M., Tino, P., 2015. Indefinite proximity learning - a review. Neural Computation 27 (10), 20392096.

Shawe-Taylor, J., Cristianini, N., 2004. Kernel Methods for Pattern Analysis and Discovery. Cambridge University Press.

Si, S., Hsieh, C., Dhillon, I. S., 2014. Memory efficient kernel approximation. In: Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014. Vol. 32 of JMLR Proceedings. JMLR.org, pp. 701–709.
URL http://jmlr.org/proceedings/papers/v32/si14.html

Smith, T. F., M.S., Waterman, Mar. 1981. Identification of common molecular subsequences. Journal of molecular biology 147 (1), 195–197.

Tsang, I. W., Kwok, J. T., Cheung, P., 2005. Core vector machines: Fast SVM training on very large data sets. Journal of Machine Learning Research 6, 363 392.

Villmann, T., Claussen, J. C., 2006. Magnification control in self-organizing maps and neural gas. Neural Computation 18 (2), 446469.

Williams, C. K. I., Seeger, M., 2000. Using the nyström method to speed up kernel machines. In: NIPS 2000. pp. 682–688.

Yang, J., Fan, L., 2013. A novel indefinite kernel dimensionality reduction algorithm: Weighted generalized indefinite kernel discriminant analysis. Neural Processing Letters, 1–13.

Zhang, K., Kwok, J. T., 2010. Clustered nyström method for large scale manifold learning and dimension reduction. IEEE Transactions on Neural Networks 21 (10), 1576–1587.

Zhang, K., Tsang, I. W., Kwok, J. T., 2008. Improved Nystrom low-rank approximation and error analysis. In: Proceedings of the 25th international conference on Machine learning. ICML '08. ACM, New York, NY, USA, pp. 1232–1239.